

# DS-GA 1008 - Deep Learning, Spring 2017

## Assignment 3

**Due: Tuesday, May 2nd, 2017 at 8:35pm**

This assignment is a “mini-project” on the subject of Generative Adversarial Networks (GANs). The team is given free choice on type of model, datasets and the specific task one finds interested to work on. The final evaluation would be grounded on the quality of the GAN generation, as well as the difficulty of the dataset/task. The novelty will be highly preferred.

---

## 1 Generative Adversarial Networks [30 credits]

1. Explain Generative Modeling.
2. Compare Generative Adversarial Networks with other Unsupervised learning approaches, such as Auto-encoders. Explain the difference.
3. Explain conditional generation using GANs, versus the vanilla unconditional version. Please briefly draw a diagram when training conditional GANs, with the condition context  $C$ , generator  $G$ , discriminator  $D$ , random vector  $z$  and output  $x$ .

[1] Goodfellow, Ian, et al. "Generative adversarial nets."

[2] Goodfellow, Ian. "NIPS 2016 Tutorial: Generative Adversarial Networks."

## 2 GAN workhouse [70 credits]

Select the model/dataset/task that is most inspirational to you, and train a GAN on it.

### 2.1 Model

[1] DCGAN. "Unsupervised representation learning with deep convolutional generative adversarial networks"

[2] WGAN. "NIPS 2016 Tutorial: Generative Adversarial Networks."

[3] EBGAN. "Energy-based generative adversarial network".

[4] LSGAN. "Least squares generative adversarial networks".

[5] BEGAN (the combination of [2] and [3]). "BEGAN: Boundary Equilibrium Generative

Adversarial Networks"

[6] ... and many many...

## 2.2 Dataset

[1] MNIST (This would serve a good starting point.)

[2] CelebA

[3] LSUN bedroom.

[4] ImageNet (definitely try taking sub-tree from WordNet).

[5] ... and many many...

Novelty comes from your own mind:

[1] Pokemon Go

[2] Google maps

[3] Social Media Stickers

[4] ... and many many...

## 2.3 Task

The GAN-related task is categorized roughly into conditional and unconditional. The vanilla GAN setting does an unconditional generation job, where it takes  $z$  and gives you an image. On the other hand, the conditional setting would be more of real value. We provide several paths digging in the conditional setting. Note this doesn't mean we dislike unconditional GANs :)

[1] Class conditional.

Conditional Generative Adversarial Nets.

[2] Factor disentanglement and style transfer.

Disentangling factors of variation in deep representation using adversarial training.

Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks.

[3] Video prediction.

Deep multi-scale video prediction beyond mean square error.

[4] Super resolution.

Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network.

[5] Semantic Segmentation using Adversarial Networks.

[6] Text-to-Image, the reverse process of image captioning.

[7] ... and many many...

## 2.4 Submission

1. Write-up.

- Answer to question 2.1
- Description of your selected model/dataset/task.

- Experiment inventory, such as what hyper-parameters you tried and what works what not.
- 50 generated photos from your best model, in a 10x5 grid. Resize your images to fit into the page.
- **If you choose to use open source code to do this assignment, please address the change you make, and describe (better with generation) what it brings you, in the inventory section.**

## 2. Model and testing script.

Send your submission (write-up, `gan_model.py`, `gan_generate.py`) to `dsga1008.spring.2k17@gmail.com` in a single zip folder named as "YourTeamName.zip" by the deadline. Include a link to the trained model file in the email. Please use the following title for your email. The testing script should give the generation from the model file (note it won't produce the exact generated images which you put in the write-up due to the random seed, don't panic).

[DS-GA-1008 YOUR\_TEAM\_NAME] Submission A3